arXiv:1704.00416v1 [q-fin.PM] 3 Apr 2017

# Sharp Target Range Strategies with Application to Dynamic Portfolio Selection

Rongju Zhang*, Nicolas Langrené†, Yu Tian‡,
Zili Zhu§, Fima Klebaner¶ and Kais Hamza‖

**Abstract**

A family of sharp target range strategies is presented for portfolio selection problems. Our proposed strategy maximizes the expected portfolio value within a target range, composed of a conservative lower target representing capital guarantee and a desired upper target representing investment goal. This strategy favorably shapes the entire probability distribution of return, as it simultaneously seeks a high expected return, cuts off downside risk, and implicitly caps volatility, skewness and other higher moments of the return distribution.

To illustrate the effectiveness of our new investment strategies, we study a multi-period portfolio selection problem with transaction cost, where the results are generated by the Least-Squares Monte-Carlo algorithm. Our numerical tests show that the presented strategy produces a better efficient frontier, a better trade-off between return and downside risk, and a wider range of possible risk profiles than classical constant relative risk aversion utility.

Finally, straightforward extensions of the sharp target range are presented, such as purely maximizing the probability of achieving the target range, adding an explicit target range for realized volatility, and defining the range bounds as excess return over a stochastic benchmark, for example, stock index or inflation rate. These practical extensions make the approach applicable to a wide array of investment funds, including pension funds, controlled-volatility funds, and index-tracking funds.

**Key words**: portfolio optimization; fund management; target strategy; performance measure; volatility target; risk measure; risk management; dynamic portfolio selection; dynamic programming; least-squares Monte Carlo; stochastic control

**JEL Classification**: G11, D81, C63, C34, **MSC Classification**: 91G10, 91G80, 91G60

*School of Mathematical Sciences, Monash University; Email: rongju.zhang@monash.edu
†Real Options and Financial Risk, Data61, CSIRO; Email: nicolas.langrene@data61.csiro.au
‡School of Mathematical Sciences, Monash University; Email: yu.tian@monash.edu
§Real Options and Financial Risk, Data61, CSIRO; Email: zili.zhu@data61.csiro.au
¶School of Mathematical Sciences, Monash University; Email: fima.klebaner@monash.edu
‖School of Mathematical Sciences, Monash University; Email: kais.hamza@monash.edu

# 1 Introduction

A crucial and long-standing problem in the theory and practice of portfolio allocation is the choice of an effective and transparent performance criteria that accommodates the trade-off between return and different measures of risk. The most famous criteria traces back to the mean-variance framework proposed by Markowitz (1952), where risk is measured by the variance of returns.

A first strand of the literature extends this mean-variance approach to incorporate higher moments, which play an important role when the return distribution is non-normal. See, for example, Lai (1991), Konno, Shirakawa, and Yamazaki (1993) incorporating a skewness component and Davies, Kat, and Lu (2009) incorporating skewness and kurtosis. The usual optimization approach is here to solve a linear objective function that accounts for higher moments through some reward/penalization weights, as a weighted sum is the easiest way to combine several, possibly conflicting objectives, see Köksalan and Wallenius (2012). These weights, designed as investment preferences for controlling risk-return trade-offs, lack however of transparency for investors, making this approach shunned in practice.

A second strand of the literature replaces variance by a better risk measure in the mean-variance framework. The most common risk measures chosen to replace variance are Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR), see Rockafellar and Uryasev (2000), Jorion (2007) and Rockafellar (2007). These measures are better than the symmetric variance at addressing downside risk, and are suitable for non-normal returns. However, the risk preference is once again accounted for by one abstract parameter, a probability threshold $p$ such that $p\%$ of the return distribution lies below $\text{VaR}_p$. In particular, the return threshold $\text{VaR}_p$ is implicit, based on the choice of $p$.

A third strand of the literature corresponds to the classical expected utility approach (von Neumann and Morgenstern (1944)), for which the effects of higher moments on investment decisions can be summed up in one risk aversion parameter, see, for example, Hodges (1998), Stutzer (2000) and Bacon (2013) for exponential, constant absolute risk aversion (CARA) utility and Zakamouline and Koekebakker (2009) for general utility. The main limit of the expected utility approach for practical purposes is the choice of utility function itself: though mathematically elegant, utility functions are abstract, impractical and the risk-return trade-off is implicit and opaque. Moreover, the researchers who did try to calibrate utility functions for practical purposes found substantial inadequacies with how people actually make decisions in practice. In particular, Tversky and Kahneman (1992)'s cumulative prospect theory has identified some practical features ignored by the classical utility theory, such as preferences relative to initial wealth, risk-seeking behavior when losing money, and overweighting of unlikely events, see Barberis (2012).

A fourth strand of literature tracks a prespecified profit target. Some of these target strategies are built upon the classical approaches discussed previously. Regarding the mean-variance approach, Franks (1992) replaces the expected return in the mean-variance framework by expected return in excess of a predefined target, and Williams (1997) maximizes the probability of beating a target return for portfolios on the efficient mean-variance frontier. Regarding the expected utility approach, Teplá (2001) maximizes an expected hyperbolic absolute risk aversion (HARA) utility under constraint of exceeding a stochastic benchmark. In general though, the most common use of predefined return targets is to maximize the probability of achieving the target. For example, Browne (1999a) and Browne (1999b) maximize the probability of exceeding a fixed target return and a stochastic benchmark respectively. Pham (2003)

solves the problem of maximizing the probability of beating a stochastic benchmark by a given percentage. Gaivoronski, Krylov, and van der Wijst (2005) propose a dynamic benchmark-tracking strategy with transaction cost, applicable to a variety of risk measures. Morton, Popova, and Ivilina (2006) maximize the probability of outperforming a benchmark with penalization of the expected downside deviation from the target. Symmetrically, one can instead minimize the probability of an undesirable outcome. Brogan and Stidham Jr. (2005) and Klebaner, Landsman, Makov, and Yao (2016) minimize the downside deviation of returns below some prespecified target (semivariance and other types of lower partial moments) subject to a given expected return constraint. Based on large deviation theory, Hata, Nagai, and Sheu (2010) and Nagai (2012) minimize the probability of underperforming a target return. As an application, Milevsky, Moore, and Young (2006) derive an optimal annuitization strategy that minimizes the probability of lifetime ruin of a retiree. Unlike mean-variance, high moment penalization, expected utility and VaR-based criteria, the explicitly specified investment return target makes this strand of literature the most likely to be understood and applied in practice by investors and traditional fund managers. However, the difficulty lies in the choice of the target level. Setting a high return target will result in a risky allocation with an unspecified propensity to lose capital. Conversely, setting the return target at too low a level will forego potential gain opportunities.

In this paper, we propose a new family of investment strategies with a prespecified range of return targets. The building block of this family is the sharp target range strategy. The idea is as follows: a return range is provided by the investor: a conservative lower target and an desired upper target. The lower target represents the level of capital guarantee. The upper target corresponds to the return level the investor wishes to achieve. The sharp target range strategy maximizes the expected return within this range, and the resulting optimal allocation decisions therefore jointly maximize the probability that the investment return lies within the target range *and* as close to the upper target as possible.

There are three main motivations behind the proposed sharp target range strategy. The first motivation traces back to the primary purpose of an objective function, which is to carve a desirable shape for the probability distribution of returns. Our target range strategy, seeking high expected return while chopping off the tails of the distribution beyond the target range, monitors the entire shape of return distribution, and the statistical moments are implicitly contained. Our numerical results show that portfolio returns are well contained within the investor's chosen range, and that our strategy dominates CRRA utility in terms of efficient frontier and trade-off between expected return and downside risk. This practical success can be attributed to the bold upside potential sacrifice above the proposed upper target: unlike classical downside risk minimization approaches, which unwisely ignore the right tail of return distribution, the proposed upper target disallows any excessive risk taking once the predefined investment goal is close to be achieved, which drastically improves downside risk management and makes the portfolio performance more predictable.

The second motivation is to ensure the ease and safety of setting the parameters of the decision criteria. Classical single-target strategies suffer from the difficulty of setting the target level, as it is very hard for one single parameter to reconcile pursuit of return while simultaneously providing downside protection. In particular, the resulting frail compromise is vulnerable to misspecification of the target: for example, setting the single target too high (while expecting both high return and downside protection) will result in excessively risky allocations with potentially extreme downside risk, in contradiction with the initial investment requirements. By contrast, our target range strategy solves this dilemma by combining one

upper target that accounts for return-seeking preference, with a lower target that independently accounts for loss-aversion preference. This separation improves the robustness to target misspecification. For example, setting the upper target too high would not result in downside risk explosion, thanks to the capital protection provided by the independent lower target.

Finally, the third motivation is to improve the communication and practical adoption of portfolio allocation techniques from academia to the industry. Performance criteria depending on abstract parameters with unforeseeable practical effects are understandably unlikely to be adopted by investors. Our proposition of two explicit targets labeled in terms of returns, with intuitive purpose (capital protection for the lower target, desired investment return for the upper target), serves this quest for a simple and practical investment criteria that provides guidance towards better portfolio allocation decisions.

To demonstrate the effectiveness of the presented range target strategy, we study a multiperiod (or dynamic) portfolio selection problem with proportional transaction cost, using the Least-Squares Monte Carlo (LSMC) algorithm presented in Zhang, Langrené, Tian, Zhu, Klebaner, and Hamza (2016). Numerical results on ex-post distributions, comparison to CRRA utility and sensitivity analysis are discussed. Moreover, the sharp target range strategy combined with the LSMC algorithm is very versatile and easy to extend to more general problems. We hereby provide three extensions. The first extension deals with pure maximization of the probability that the final portfolio value ends within the target range, without further attempt to pursue a high expected return. As will be discussed in 5, this objective is useful whenever avoiding ruin is more important than earning high returns, for example for pension schemes, retirement funds and life-cycle management. The second extension incorporates a second, explicit target range on the realized volatility. This allows for a much more precise control over realized volatility, which is very useful for funds that explicitly target a low volatility or low tracking error. Our results show that a minimum of 3%-4% one-year monthly realized volatility can be achieved for a portfolio of cash, bond and stocks. The third extension concerns relative returns: the bounds of the sharp target range can be defined in terms of excess return over a stochastic benchmark, such as stock index, interest rate or inflation rate.

The outline of the paper is as follows. Section 2 describes the dynamic portfolio selection problem and the sharp target range strategy. Section 3 describes the implementation of the LSMC algorithm. Section 4 discusses our numerical findings for the sharp target range strategy, including out-of-sample distributions, benchmarking with CRRA utility and sensitivity analysis. Section 5 extends the sharp target range strategy to the probability-maximizing, flat target range strategy and numerical out-of-sample distribution results are discussed. Section 6 extends the sharp target range strategy to incorporate a realized volatility target component and discusses numerical results on out-of-sample distributions and sensitivity analysis. Section 7 extends our sharp target range strategy from fixed target to stochastic target and the out-of-sample distributions are presented. Section 8 concludes the paper.

## 2 Sharp target range strategy

In this section, we mathematically define the proposed sharp target range strategy and describe the dynamic portfolio optimization problem. Consider a multiperiod portfolio optimization problem over a finite time horizon $T$. Let $\mathcal{T} = \{1, \dots, T-1\}$ be the discrete set of portfolio rebalancing times. Let $N$ be the number of possible risky assets in the portfolio and $\alpha_t = \left(\alpha_t^i\right)_{i=1,\dots,N}$ be the portfolio weights in

each risky asset at time $t \in \mathcal{T}$. Let $\mathcal{C} \subseteq \mathbb{R}^N$ be the set of admissible strategies for all the rebalancing times $t \in \mathcal{T}$. Let $W_t$ be the portfolio value at time $t \in \mathcal{T}$. For each rebalancing time $t \in \mathcal{T}$, the portfolio weights are rebalanced based on the available information $\mathcal{F}_t$ at time $t$ (which contains the asset prices, the portfolio value $W_t$, and possibly additional economic predictors $Z_t$, such as oil price, gold price, interest rates, etc.). Our dynamic portfolio selection problem reads

$$V_t = \sup_{\{\alpha_s \in \mathcal{C}\}, s \in \mathcal{T} \cap [t,T)} \mathbb{E}[f(W_T)|\mathcal{F}_t], \ 0 \leq t \leq T \tag{2.1}$$

where the objective function depends on the final wealth $W_T$ through the function $f$. The wealth $W$ is an endogenous state variable which depends on all the rebalancing decisions $\alpha$; in particular transaction costs are taken into account. As for the definition of $f$, we propose the following parametric shape:

$$f(w) = (w - L_W)\mathbb{1}\{L_W \leq w \leq U_W\} \tag{2.2}$$

Equation (2.2) depends on two parameters: a conservative lower target $L_W \in \mathbb{R}$ and a desired upper target $U_W \in \mathbb{R}$. Using the shape (2.2), the objective function (2.1) maximizes expected final wealth within the interval $[L_W, U_W]$, while the values outside this interval are penalized down to zero. This strategy implicitly combines two goals: that the portfolio value ends within the chosen target range $[L_W, U_W]$, *and* close to the upper target $U_W$. This parametric shape has several advantages, some of which will be discussed in more details in the numerical section 4 of the paper:

- It is simple yet allows for various types of risk preferences. For instance, a tight tracking can be achieved by setting a tight interval, while targeting high returns is possible by setting a high upper target.
- Accounting for the risk preferences of the investor is easy and transparent, as it depends on only two, easily interpretable parameters.
- It caps volatility and other higher moments better than the mean-variance approach and the traditional moments-constrained approaches, as it tightens the entire distribution of terminal portfolio value.
- As the objective function (2.1) does not have any explicit chance constraint or expectation constraint, solving the multiperiod allocation problem (2.1) is as easy as performing any classical expected utility maximization.

Equation (2.1) reveals that $f$ can be mathematically interpreted as a non-classical truncated linear "utility" function. It is worth noting that the upper target $U_W$ acts as a limit to the upside potential, which is at odds with classical utility theory. The reason for this choice is that this upside sacrifice is precisely what makes the sharp target range (2.2) effective at eliminating downside risk and reducing volatility and higher moments of terminal wealth. Above all, the choice (2.2) makes realized returns much more predictable, which in many contexts is more important than keeping the possibility of rare windfall returns at the cost of higher downside risk.

Throughout this paper, we normalize the portfolio value $W$ and the bounds $[L_W, U_W]$ by the initial wealth $W_0$. Indeed, formula (2.2) shows that $f(W; L_W, U_W) = W_0 \times f(\frac{W}{W_0}; \frac{L_W}{W_0}, \frac{U_W}{W_0})$, so that we can assume without loss of generality that $W_0 = 1$ and set the bound $L_W$ and $U_W$ in the vicinity of 1, accounting for the investor's risk profile. Figure 2.1 illustrates the target range (2.2) with $L_W = 1.0$ and $U_W = 1.2$.

Figure 2.1: Sharp target range

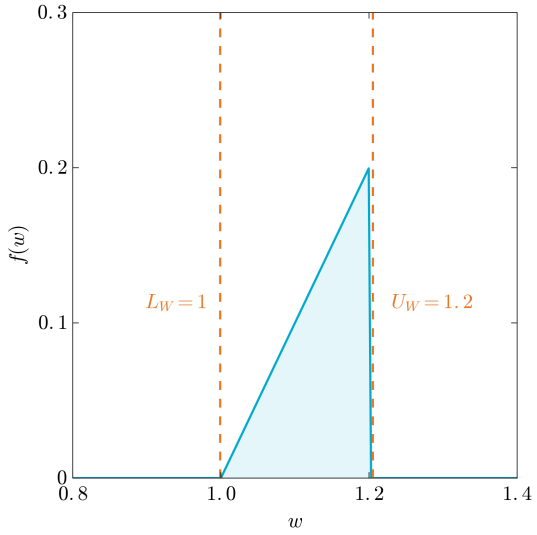


# 3 Algorithm

In this section, we formulate our portfolio selection problem as a dynamic programming problem and describe how to solve it using the LSMC algorithm. Subsection 3.1 discusses how to solve our portfolio allocation problem (2.1) using the Least-Squares Monte Carlo (LSMC) algorithm. The truncated linearity of target range (2.2) allows for simple parametric dynamic programming estimates, making the LSMC algorithm stable, robust and fast. Subsection 3.2 provides additional implementation details, such as how to deal with the endogenous wealth variable, how to refine the allocation grid, and whether to use value function iteration or policy function iteration.

## 3.1 Overview

In this subsection, we discuss how to solve our problem (2.1). Our goal is to maximize the expected terminal portfolio value within a target range. The optimization problem reads

$$V_t = \sup_{\{\alpha_s \in \mathcal{C}\}, s \in \mathcal{T} \cap [t,T)} \mathbb{E}[(W_T - L_W)\mathbb{1}\{L_W \leq W_T \leq U_W\} | X_t, \alpha_{t^-}], \ t \in \mathcal{T} \tag{3.1}$$

where the Markovian vector $X_t$ contains all the state variables other than the portfolio weights $\alpha_t$ (current asset prices, current portfolio value, additional economic variables, etc.). To solve this dynamic problem, we rely on the application of the LSMC algorithm to portfolio allocation developed in Zhang et al. (2016). The algorithm is composed of two main parts. The first one corresponds to the forward Monte Carlo simulations of all the state variables. The second part corresponds to the backward dynamic programming loop, where the conditional expectations (continuation values) are estimated by least-squares regression. More precisely, the dynamic programming equation reads

$$V_t = V_t(X_t, \alpha_{t^-}) = \max_{a^j \in \mathcal{C}_d} \mathbb{E}\left[ f(W_T) \,\middle|\, X_t, \alpha_t = a^j \right] \ , t \in \mathcal{T} \tag{3.2}$$

where, for implementation purposes, we assume a discrete set $\mathcal{C}_{\mathrm{d}} := \{a^j\}_{j=1,\dots,\#(\mathcal{C}_{\mathrm{d}})} \subseteq \mathcal{C}$ of admissible strategies.

In the LSMC algorithm, the continuation values $CV_t^j(x) := \mathbb{E}\left[V_{t+1} \left| x, \alpha_t = a^j \right.\right] = \mathbb{E}\left[f(W_T) \left| x, \alpha_t = a^j \right.\right]$ are estimated by least-squares regression: for each $a^j \in \mathcal{C}_{\mathrm{d}}$, the Monte Carlo sample $\{f(W_T^{j,m})\}_{m=1,\dots,M}$ (policy function iteration) or the sample $\{V_{t+1}^{j,m}\}_{m=1,\dots,M}$ (value function iteration, see Subsection 3.2) is regressed with respect to a linear combination of some basis functions of the input sample $\{X_t^m\}_{m=1,\dots,M}$, where $M$ is the number of Monte Carlo simulations.

In our case, the objective function $f$ has a truncated linear shape (equation (2.2)), which would be difficult to fit in a stable manner with a classical global basis (polynomial basis for example). The difficulty lies with the domain where $f(W_T)$ is equal to zero. To deal with this problem, we adapt the classical Tobit model Tobin (1958)'s idea on censored regression. First, we approximate $W_T$ as a linear combination of some basis functions $\psi(X_t^j)$ of $X_t^j := (X_t, a^j)$:

$$W_T = \beta \cdot \psi(X_t^j) + \sigma\varepsilon, \ \varepsilon \sim \text{i.i.d. } \mathcal{N}(0,1), \tag{3.3}$$

where $\varepsilon$ follows a standard normal distribution with zero mean and unit variance. Let $\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(\frac{x^2}{2}\right)$ represent the standard normal probability density function, and $\Phi(x) = \int_{-\infty}^{x} \phi(x)dx$ represent the standard normal cumulative distribution function. The random variable $(W_T - L_w)\,\mathbb{1}\{L_w \le W_T \le U_w\}$ follows a truncated normal distribution and its expectation can be easily derived:

$$\begin{aligned}
\mathbb{E}[f(W_T)|X_t^j] &= \mathbb{E}\left[(W_T - L_w)\,\mathbb{1}\{L_w \le W_T \le U_w\}|X_t^j\right] \\
&= \mathbb{E}\left[\left(\beta \cdot \psi(X_t^j) - L_w + \sigma\varepsilon\right)\mathbb{1}\left\{L_w - \beta \cdot \psi(X_t^j) \le \sigma\varepsilon \le U_w - \beta \cdot \psi(X_t^j)\right\}|X_t^j\right] \\
&= \left[\beta \cdot \psi(X_t^j) - L_w\right]\left\{\Phi\left(\frac{U_w - \beta \cdot \psi(X_t^j)}{\sigma}\right) - \Phi\left(\frac{L_w - \beta \cdot \psi(X_t^j)}{\sigma}\right)\right\} \\
&\quad - \sigma\left\{\phi\left(\frac{U_w - \beta \cdot \psi(X_t^j)}{\sigma}\right) - \phi\left(\frac{L_w - \beta \cdot \psi(X_t^j)}{\sigma}\right)\right\}.
\end{aligned} \tag{3.4}$$

To sum up, one can first estimate the coefficients $\beta$ and $\sigma$ by ordinary least-squares regression (equation (3.3)). Then, thanks to the truncated linear shape of (2.2), the conditional expectations in the dynamic programming equations (3.2) can be estimated by the elegant closed-form formula (3.4). This two-stage approach makes the LSMC algorithm stable and fast. In particular, the truncated linear shape (2.2) works much better than classical concave utility functions within the LSMC, for which a classical direct approximation $\mathbb{E}[U(W_T)|X_t^j] \approx \beta \cdot \psi(X_t^j)$ can lead to a large cumulative Monte Carlo error when risk-aversion is high (strong concavity), see Van Binsbergen and Brandt (2007), Garlappi and Skoulakis (2009) and Zhang et al. (2016). Our objective function (2.2) avoids this non-linearity problem altogether, and the LSMC algorithm converges very quickly when using the parametric approach (3.3)-(3.4).

More generally, the approach adopted here, based on the Tobit approach (linear approximation (3.3) + corrections due to truncation (3.4)) can also be adapted to the case when residuals are significantly non-Gaussian (this would simply modify the correction terms (3.4)). There is no restriction on the choice of return distribution, nor on the estimation methods to use for the residual distributions if needed (empirical distribution, kernel estimation, mixture normal, etc.). Nevertheless, for low-frequency trading such as monthly returns with monthly rebalancing considered in our numerical section 4, the

normality assumption is reasonable. For this reason, and as our main focus is on the analysis of the new investment strategy (2.2), we henceforth assume normality of residuals and use the approach (3.3)-(3.4) for our implementation.

## 3.2 Implementation details

In this Subsection, we provide additional information on a few key implementation details. For a comprehensive description of the application of the LSMC algorithm to portfolio allocation problems with intermediate costs, we refer to Zhang et al. (2016). The reader more interested in the numerical performance of the sharp target range strategy than in the detailed implementation of the algorithm may skip this subsection and move to the section on numerical results 4.

**Control randomization**

As previously mentioned, the LSMC algorithm is composed of two main parts: a forward simulation loop and a backward dynamic programming loop. Forward simulation consists in simulating a large sample of all the stochastic risk factors. Exogenous risk factors such as asset returns are easy to simulate. However, the portfolio value, which depends on the rebalancing decisions due to transaction costs, cannot be directly simulated. To deal with this endogenous risk factor, we use the so-called control randomization technique (Kharroubi, Langrené, and Pham 2014). First, we model decisions as additional dummy random risk factors in the forward loop. Then, during the backward loop, the dummy decisions are made part of the conditioning information, i.e. they are modeled as the current portfolio allocation state, a piece a information required to compute the incurred transaction costs. This provides continuation value estimates for each possible decision $a^j \in \mathcal{C}_\mathrm{d}$, so that the dynamic programming equation (3.2) can be solved for each simulation path.

After this short summary, we provide below further details about the backward loop. At time $T-1$, for each possible decision $a^j \in \mathcal{C}_\mathrm{d}$, we recompute the simulated endogenous wealth variable forward for each path, assuming $a^j$ as the current allocation just before the $T-1$ rebalancing. We then follow equation (3.3) and use ordinary least-squares regression to approximate the recomputed terminal wealth as a linear combination of some basis of all the state variables. This yields the parameter estimates

$$\hat{\beta}^j_{T-1} := \arg\min_{\beta\in\mathbb{R}^p} \frac{1}{M}\sum_{m=1}^{M}\left(W_T^m - \psi(X_{T-1}^{j,m})\cdot\beta\right)^2,$$

$$\hat{\sigma}^j_{T-1} := \sqrt{\frac{1}{M-p}\sum_{m=1}^{M}\left(W_T^m - \psi(X_{T-1}^{j,m})\cdot\hat{\beta}^j_{T-1}\right)^2},$$

where $p$ denotes the number of least-squares coefficients. We then simply use these two estimates to compute the continuation value estimate (3.4) for each Monte Carlo path, and pick the optimal decision that maximizes the continuation value (dynamic programming equation (3.2)).

We proceed similarly at time $t = T-2, ..., 1$. When recomputing the endogenous wealth forward from its current value for each Monte Carlo path and each current possible allocation $a^j \in \mathcal{C}_\mathrm{d}$, we use the previously estimated parameters $(\hat{\beta}^l_{t+1}, \hat{\sigma}^l_{t+1}), ..., (\hat{\beta}^l_{T-1}, \hat{\sigma}^l_{T-1})$, $l|a^l \in \mathcal{C}_\mathrm{d}$ to recompute the future optimal decisions that maximize the estimated continuation values (3.4) (the optimal future policy functions do

not change, but the actual future decisions for each path can change because the current optimal decisions affect the future wealth level). We record the current estimates $\hat{\beta}_t^j$ and $\hat{\sigma}_t^j$ and keep going until time 0. At time 0, the procedure is easier. As the initial information is observable, there is no need for estimation of conditional expectation. The initial optimal decision is chosen simply by maximizing the continuation value (which at time 0 is a simple expectation) over all the possible decisions $a^j \in \mathcal{C}_\mathrm{d}$.

**Local interpolation**

Regarding the maximizations, we are so far using a simple exhaustive search over a fixed grid $\mathcal{C}_\mathrm{d}$ of potential portfolio allocations. While simple and straightforward, this method can become computationally expensive when the grid mesh is fine and the number of assets is high. To alleviate this curse of dimensionality with respect to the action space, we propose to combine an exhaustive search over a coarse grid $\mathcal{C}_\mathrm{d}$ with adaptive local interpolation over the allocation space.

For each $t \in \mathcal{T}$, we first build a coarse grid $\mathcal{C}_\mathrm{d} \subseteq \mathcal{C}$ and estimate $\hat{a}_t^* = \max_{a^j \in \mathcal{C}_\mathrm{d}} \mathbb{E}[f(W_T)|X_t^j]$. Then, we extract the local hypercube $\mathcal{L}(\hat{a}_t^*) \subseteq \mathcal{C}_\mathrm{d}$ around $\hat{a}_t^*$, and interpolate $\mathbb{E}[f(W_T)|X_t^j]$ parametrically inside the interior of $\mathcal{L}(\hat{a}_t^*)$, denoted $\overline{\mathcal{L}}(\hat{a}_t^*)$. Finally, we perform the parametric maximization

$$\alpha_t^* \quad = \quad \arg \sup_{a \in \overline{\mathcal{L}}(\hat{a}_t^*)} \mathbb{E}[f(W_T)|X_t, a].$$

In our experience, even a very coarse discrete portfolio allocation grid $\mathcal{C}_\mathrm{d}$ can achieve results similar to a fine grid, for a fraction of the computational cost, when local interpolation is used. We refer to the numerical tests in Zhang et al. (2016).

**Policy function iteration**

In view of the dynamic programming equation (3.2), there are two ways to implement the LSMC algorithm. The value function iteration (VFI, Carriere (1996), Tsitsiklis and Van Roy (2001)) approximates the value function by regression, i.e. uses the regression estimate of $\mathbb{E}[\,V_{t+1}|\,X_t^j]$ as the estimated continuation value. By contrast, the policy function iteration (PFI, Longstaff and Schwartz (2001)) discards the regression estimate once the optimal policy has been computed, and only keeps track of the optimal policies until the end of the investment horizon. Subsection 3.1 shows that we adopt the PFI approach here. As the portfolio value $W$ is endogenous, PFI requires $W_T$ to be recomputed for every $t \in \mathcal{T}$ in the backward loop, after which $\mathbb{E}[\,f(W_T)|\,X_t^j]$ can be estimated by regression to compute the current optimal decisions.

In practice, this means that the computational complexity of PFI grows quadratically in the number of rebalancing $\#\mathcal{T}$, while it grows only linearly for VFI. However, the advantage is that PFI is always more stable than VFI, as regression errors occur only through the estimated decisions $\alpha_t^*$, while VFI compounds all the regression errors on the continuation value estimates themselves. When there exists a natural, suitable guess for the regression basis, the regression errors can be limited, and VFI and PFI produce similar results (Garlappi and Skoulakis (2009), Zhang et al. (2016)). However, in the general case, it is hard to exhibit a suitable global regression basis (which is notably the case in this paper, due to the downward jump of the objective function at the upper target $U_w$), in which case PFI produces better results, as it is more robust to the choice of basis.

# 4 Numerical experiments

In this section, we illustrate numerically that the proposed sharp target range strategy (2.2) is able to consistently meet the investor's range objective. Table 4.1 summarizes the asset classes and state variables used for our numerical experiments. We consider a portfolio invested in five assets: risk-free cash, U.S. bonds (AGG), U.S. shares (SPY), international shares (IFA) and emerging market shares (EEM), the other assets listed in Table 4.1 being used as potential return predictors. The intermediate transaction cost is assumed to be a fixed proportion of the portfolio turnover, set to 0.1%. The annual interest rate on the cash component is set to be 2%.

A first-order vector autoregression (VAR) model is calibrated to the monthly log-returns of Table 4.1's assets from September 2003 to March 2016. By bootstrapping the residuals, 10000 Monte-Carlo paths are generated as simulation paths for a twelve-month forward time period. 10000 paths is more than enough for numerical stability of our algorithm; in fact, the stability of the two-stage regression approach introduced in subsection 3.1 is such that as little as 1000 paths already produce stable results.

Table 4.1: Assets

| Assets | Underlying | Data source |
|---|---|---|
| U.S. Bonds | AGG (ETF) | Yahoo Finance |
| U.S. Shares | SPY (ETF) | Yahoo Finance |
| International Shares | IFA (ETF) | Yahoo Finance |
| Emerging Market Shares | EEM (ETF) | Yahoo Finance |
| Japanese shares | Nikkei225 | Yahoo Finance |
| U.K. shares | FTSE100 | Yahoo Finance |
| Australian shares | ASX200 | Yahoo Finance |
| Gold | Spot Price | World Gold Council |
| Crude Oil | Spot Price | U.S. Energy Info. Admin. |
| U.S. Dollar | USD Index | Federal Reserve |
| Japanese Yen | JPYUSD | Federal Reserve |
| Euro | USDEUR | Federal Reserve |
| Australian Dollar | USDAUD | Federal Reserve |

For the least-squares regression, firstly, a principal component analysis is performed on the thirteen exogenous state variables and only the components that capture more than 95% variation are used to form the underlying bases. Afterwards, a simple multivariate polynomial basis function is used in the linear least-squares regression (3.3). We choose the polynomial order by observing the plots of the value functions with respect to the principal components at time $T-1$. These plots indicate an almost linear but slightly curved (quadratic) relationship. We therefore choose polynomials of order-two (including intercept) as regression bases. By the orthogonality of the principal components, interacted terms of the polynomial bases do not need to be included, which improves the computational efficiency.

At each rebalancing time, and for every Monte Carlo path, we search for the best portfolio allocation over a discrete allocation grid with 20% mesh. This first estimate is then improved using the adaptive

local interpolation described in Subsection 3.2.

## 4.1 Out-of-sample distribution

Figure 4.1 provides some out-of-sample distributions for the value of the investment portfolio at the end of a twelve-month horizon. We recall that the portfolio value $W$ and the bounds $[L_w, U_w]$ are scaled by the initial wealth, so that without loss of generality $W_0 = 1$. The lower target $L_w$ is set to the initial wealth level 1.00, a natural choice representing investors' preference for capital protection. Four different upper targets $U_w$ are tested: 1.05, 1.10, 1.20 and 1.30. $U_w$ can be viewed as representing the 'appetite' for investment returns. Several comments can be made about the shape of the terminal wealth distribution produced by the sharp target range strategy.

Figure 4.1: Target range strategy: out-of-sample distribution of final portfolio value

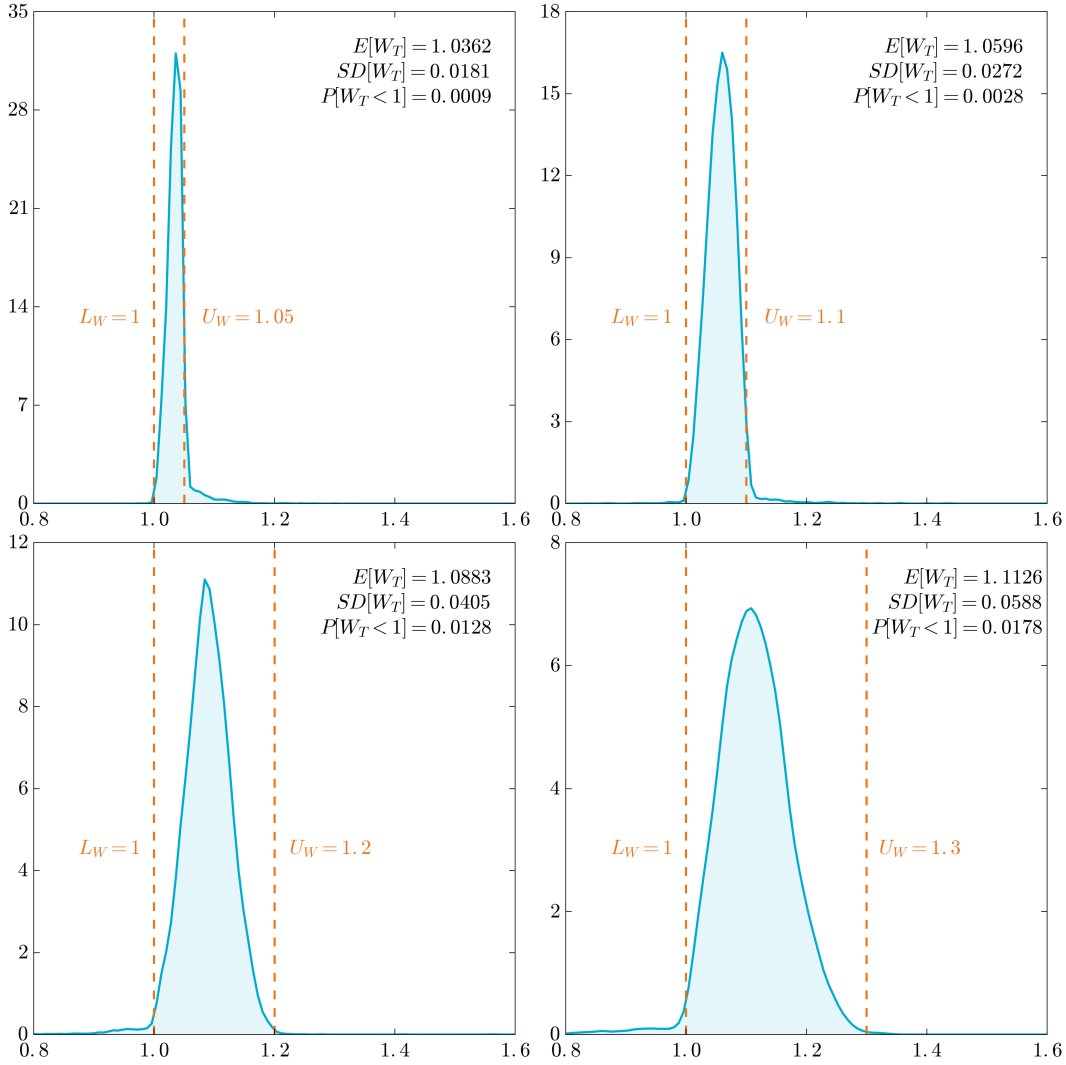


Firstly, the most striking observation is that our target range strategy (2.1) is successful at containing the wealth distribution within the predefined target range. The fact that the proposed objective function

(2.2) can successfully replicate the type of results that could be expected from a more involved chance-constrained optimization is the main finding of the paper, and the main motivation behind the parametric shape (2.2) in the first place.

Secondly, the wealth distribution mimics the goal shape as presented in Figure 2.1. In particular, setting the upper target $U_W$ to a higher level produces a higher expected final wealth with higher standard deviation and greater downside risk (as measured by the probability of losing capital). For the tails beyond the target range, the two low levels $U_W = 1.05$ and $U_W = 1.10$ produce a fatter right tail, while the two higher levels $U_W = 1.20$ and $U_W = 1.30$ produce a fatter left tail, which is consistent with the fact that the greater $U_W$, the greater the willingness to take risk so as to end as close as possible to the upper target. This observation demonstrates the capability of the sharp target range strategy to represent different risk appetites.

An interesting quantity to monitor is the ratio $\mathcal{R} := (\mathbb{E}[W_T] - L_W)/(U_W - L_W)$ which measures the location of the expected performance $\mathbb{E}[W_T]$ within the target range: $\mathcal{R} = 0\%$ means $\mathbb{E}[W_T] = L_W$, while at the opposite $\mathcal{R} = 100\%$ means $\mathbb{E}[W_T] = U_W$. In our experiments from Figure 4.1, $\mathcal{R}$ is a decreasing function of $U_W$, from $\mathcal{R} = 72\%$ for $U_W = 1.05$ down to $\mathcal{R} = 38\%$ for $U_W = 1.30$. This illustrates the natural fact that the higher the desired upper target, the harder it is to achieve it.

## 4.2 Outperforming classical utility

In this subsection, we show that the sharp target range strategy (2.2) always outperforms classical utility functions such as CRRA utility (constant relative risk aversion $u(x) = x^{1-\gamma}/(1-\gamma)$). The CRRA utility and other classical utility functions (CARA, HARA, etc.) have two main problems when it comes to solving portfolio selection problems in practice:

1. Their risk aversion parameters are abstract and difficult to set for practical investment decisions.
2. For multiperiod problems, a simulation-and-regression approach generates large numerical errors when the utility function is highly nonlinear (high risk aversion), as noted in Van Binsbergen and Brandt (2007), Garlappi and Skoulakis (2009) and Zhang et al. (2016).

By design, the alternative target range "utility" (2.2) proposed in this paper does not suffer from these problems. It is therefore interesting to perform a numerical comparison of the performance of the portfolio allocations produced by the target range (2.2) and a classical utility such as CRRA, using the same algorithm described in Section 3 with enough Monte Carlo paths to ensure convergence for the less stable CRRA utility.

Figure 4.2 compares the portfolio value distributions generated by the sharp target range strategy and the CRRA utility. Our main finding is that for each risk aversion level $\gamma$ of the CRRA utility, there exists a target range $[L_W, U_W]$ such that the sharp target range strategy produces a higher expectation, lower standard deviation and lower downside risk. As an illustration of this fact, Figure 4.2 provides two examples on how the sharp target range strategies $[L_W, U_W] = [0.9, \infty]$ and $[L_W, U_W] = [0.93, 1.53]$ outperform the CRRA utility with $\gamma = 2$ and $\gamma = 10$, respectively.

As previously discussed in Subsection 3.1, the CRRA utility can suffer from large numerical error in the LSMC algorithm when the risk-aversion parameter is very high (high nonlinearity), while the sharp

target range does not suffer from this problem at all (only a linear approximation is required, which feeds a closed-form conditional expectation estimate).

In order to allow $\gamma$ to be greater than 10 while ensuring convergence within a manageable running time, we now reduce the investment horizon to a three-month period. This allows us to obtain stable CRRA results for up to $\gamma = 30$. Figure 4.3 provides the efficient frontier obtained with range targeting (for different $U_W$ levels) and CRRA (for different $\gamma$ levels). One can see that the target range, truncated linear "utility" (2.2) always dominates CRRA utility: for any given level of standard deviation and any given level of downside risk, the sharp target range strategy achieves a higher expected final wealth than CRRA. Moreover, the target range strategy is much more flexible: in particular, it is naturally able to reach very low levels of standard deviation and downside risk (by setting $U_W$ close to $L_W$), which is hardly reachable by classical nonlinear utility functions. This is the reason why the CRRA efficient frontier is much shorter than the target range one on Figure 4.3.

Figure 4.2: Comparison with CRRA: out-of-sample distributions of final portfolio value

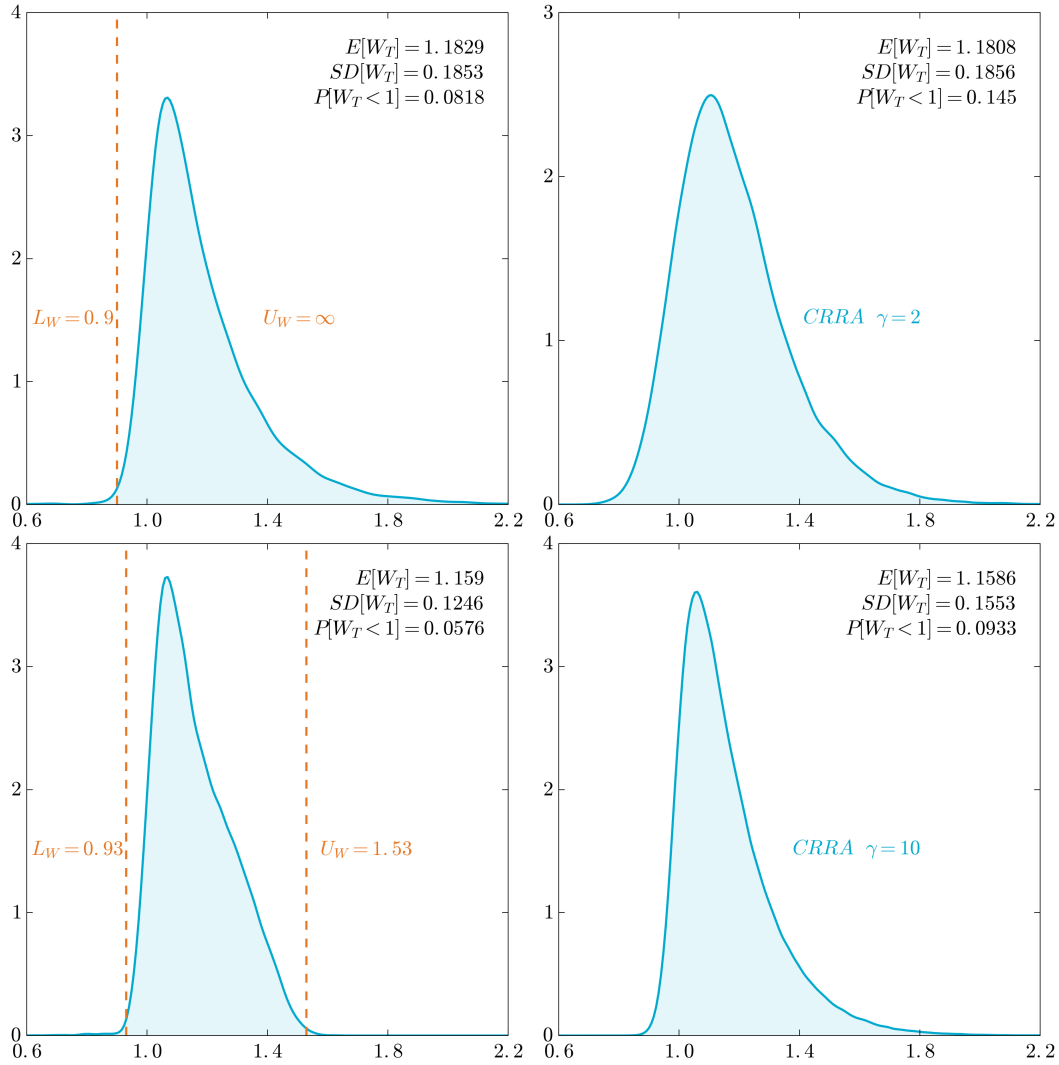


A theoretical proof of the higher efficiency of sharp target range strategies over classical utility strategies

would be desirable to corroborate our numerical findings. However, based on the difficulty to obtain an explicit optimal allocation for a single trading period with a simpler downside risk minimization objective (Klebaner, Landsman, Makov, and Yao (2016)), obtaining such a proof for sharp target range in a multiperiod, multi-asset setting might be out of reach. We leave it for further research.

Figure 4.3: Comparison with CRRA: efficient frontiers

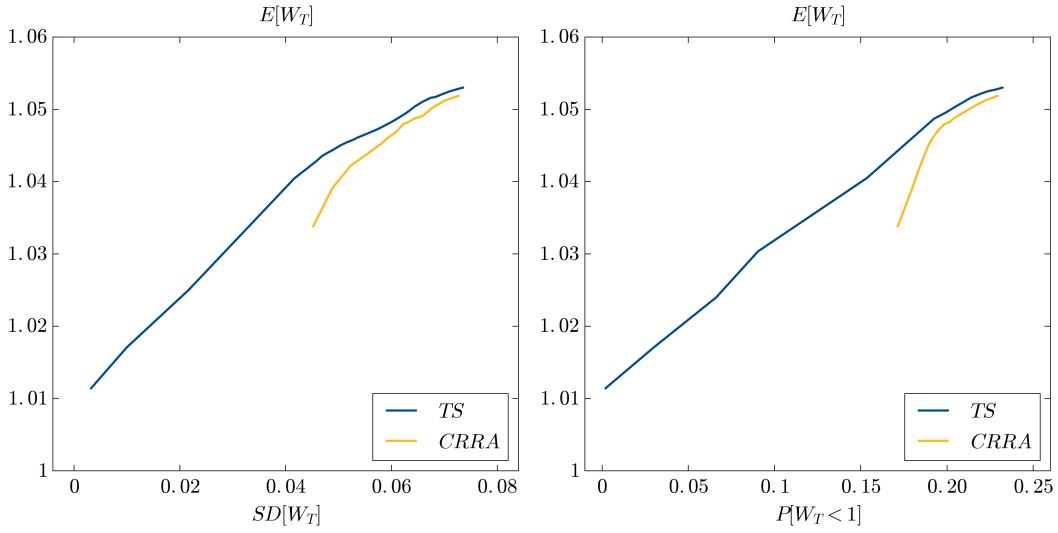


## 4.3 Sensitivity analysis

In this subsection, we analyze the sensitivities with respect to the bounds of the target range. Figure 4.4 shows how the expected terminal wealth (first row), the standard deviation of terminal wealth (second row) and the downside risk (probability of losing money, third row) are affected by a change in the upper bound $U_W$ (left column) and by a change in the lower bound $L_W$ (right column).

The main observation is that the distribution of final wealth is much more sensitive to the upper level $U_W$ than to the lower level $L_W$. The left column of Figure 4.4 shows how an increase in $U_W$ increases at the same time the expectation, standard deviation and downside risk of the final portfolio value, though a plateau is reached around 1.5 for downside risk and around 1.8 for expectation.

There is little effect of $L_W$ on the expected terminal wealth when $U_W$ is low. When $U_W$ is large, an increase in $L_W$ slightly decreases the expected terminal wealth. This is due to the additional risk required at the beginning of the trading period to grow the portfolio value from $W_0 = 1$ to the lower target $L_W$, which also expands the left tails and consequently decreases the final expectation. For the same rationale, the standard deviation and the downside risk increase when $L_w$ is increased above 1. Interestingly, the standard deviation and downside risk also increase when $L_w$ is decreased below 1 (their minimum is reached for $L_w = 1$). This time, the reason is that the penalty for falling below $L_w$ is more remote when $L_w$ is lower than 1, which induces a riskier initial portfolio.

As a conclusion, our observations indicate that $L_W = 1$ is the safest choice for the lower bound of the target interval, from which the upper bound $U_W$ can be set according to the risk preference and return requirement of the investor.

Figure 4.4: Sensitivity analysis

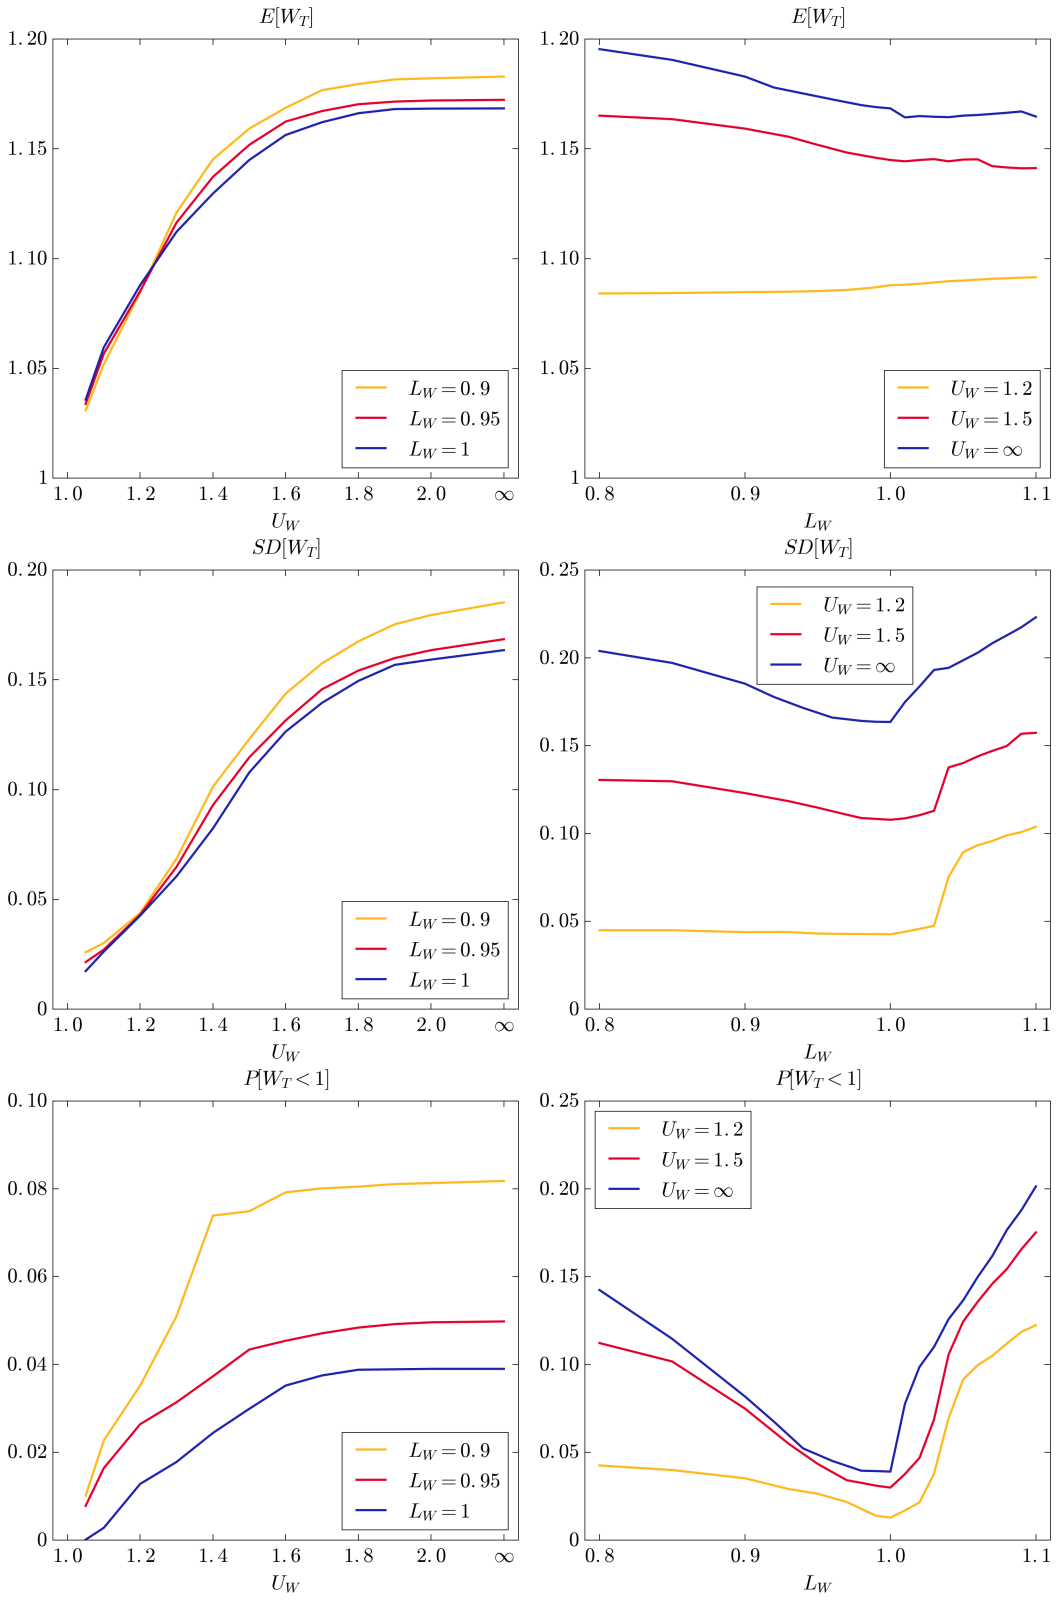

# 5 Extension 1: flat target range strategy

The last three sections will be devoted to extensions of the target range strategy (2.1), so as to illustrate the scope of the family of target range strategies, and hint at how problems involving dynamic financial decision can be addressed with the suitable type of target range.

Our first extension corresponds to flat target range strategy, in other words pure maximization of the probability of ending within the target range. Our sharp target range strategy (2.1) was tilted towards the upper return target, so as to provide a reasonable balance between return seeking and downside risk protection suitable for most classical portfolio allocation problems. Yet, there exists some other types of portfolio allocation problems for which the ability to remain solvent over time prevails over the appetite for high returns. This includes insurance-related problems such as lifecycle portfolio choice. For such problems, we propose a simple flat target range strategy

$$f(w) = \mathbb{1}\{L_w \leq w \leq U_w\}. \tag{5.1}$$

Figure 5.1 illustrates equation (5.1) with $[L_w, U_w] = [1.0, 1.2]$.

Figure 5.1: Flat target range (probability maximizing)

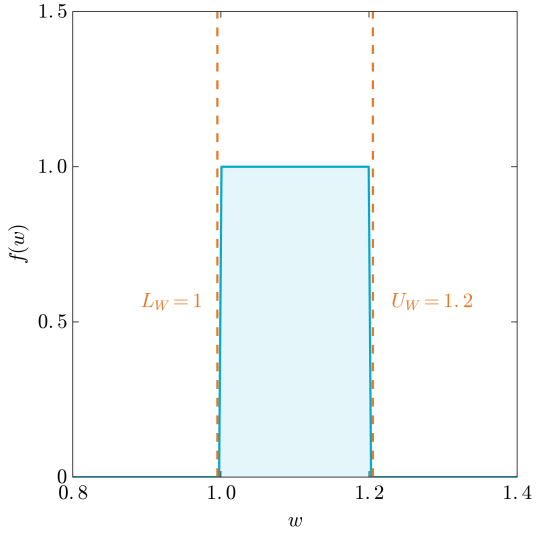


Note that when $U_w$ is set to $+\infty$, the probability maximizing strategies (5.1) and VaR minimizations produce the same range of outcomes, the difference being a fixed, absolute cut-off level for the former and an implicit, relative cut-off level for the latter. In particular, the flat target range strategy minimizes the probability of being below a particular level of 'loss', while the VaR procedure minimizes a particular loss quantile. Compared to VaR minimization, we argue that the probability-maximizing strategy provides more flexibility for investors to devise their risk preferences, as the lower return target $L_w$ is an explicit input from the investor, and the option to fix an upper target $U_w$ broadens the range of possible risk profiles.

## 5.1 Solution

The optimization problem for the flat target range strategy becomes

$$V_t = \sup_{\{\alpha_s \in \mathcal{C}\}, s \in [t, T-1]} \mathbb{E}[\mathbb{1}\{L_W \leq W_T \leq U_W\} | X_t^j, \alpha_t] = \sup_{\{\alpha_s \in \mathcal{C}\}, s \in [t, T-1]} \mathbb{P}[L_W \leq W_T \leq U_W | X_t^j, \alpha_t]. \quad (5.2)$$

The same procedure described in Section (3) can be used to solve this problem (5.2), the only difference being the computation of the continuation values. As in Section 3, we first perform a linear approximation at time $t$ of the final portfolio value $W_T$ as a function of all the state variables $X_t^j$ for each possible allocation $a^j \in \mathcal{C}_{\mathrm{d}}$:

$$W_T = \beta \cdot \psi(X_t^j) + \sigma\varepsilon, \ \varepsilon \sim \text{i.i.d. } \mathcal{N}(0,1),$$

where $\varepsilon$ follows a standard normal distribution. From there the new continuation values for flat target range are given by

$$\begin{aligned}
\mathbb{P}[L_W \leq W_T \leq U_W | X_t^j] &= \mathbb{P}[L_W \leq \beta \cdot \psi(X_t^j) + \sigma\varepsilon \leq U_W | X_t^j] \\
&= \mathbb{P}\left[\frac{L_W - \beta \cdot \psi(X_t^j)}{\sigma} \leq \varepsilon \leq \frac{U_W - \beta \cdot \psi(X_t^j)}{\sigma} \mid X_t^j\right] \\
&= \mathbb{P}\left[\varepsilon \leq \frac{U_W - \beta \cdot \psi(X_t^j)}{\sigma} \mid X_t^j\right] - \mathbb{P}\left[\varepsilon \leq \frac{L_W - \beta \cdot \psi(X_t^j)}{\sigma} \mid X_t^j\right] \\
&= \Phi\left(\frac{U_W - \beta \cdot \psi(X_t^j)}{\sigma}\right) - \Phi\left(\frac{L_W - \beta \cdot \psi(X_t^j)}{\sigma}\right).
\end{aligned}$$

We can then follow the same procedure described in Section 3 to solve the dynamic portfolio selection problem (5.2).

## 5.2 Out-of-sample distribution

Following the same numerical setting as in Section 4, we show on Figure 5.2 some out-of-sample distributions of final portfolio value produced by the probability maximization strategy .

The main observation is that, as expected, the probability of ending outside the predefined range $[L_W, U_W]$ is much smaller than the one resulting from the sharp target strategy (2.2) (refer to Figure 4.1 for comparison). This is the main strength of the flat target range (5.1): downside risk is kept to a minimum. The price to pay for this safety is the lost ability to generate returns close to the upper target $U_W$, as shown by the lower expected returns on Figure 5.2 compared to Figure 4.1.

Figure 5.2: Probability maximizing strategy: out-of-sample distribution

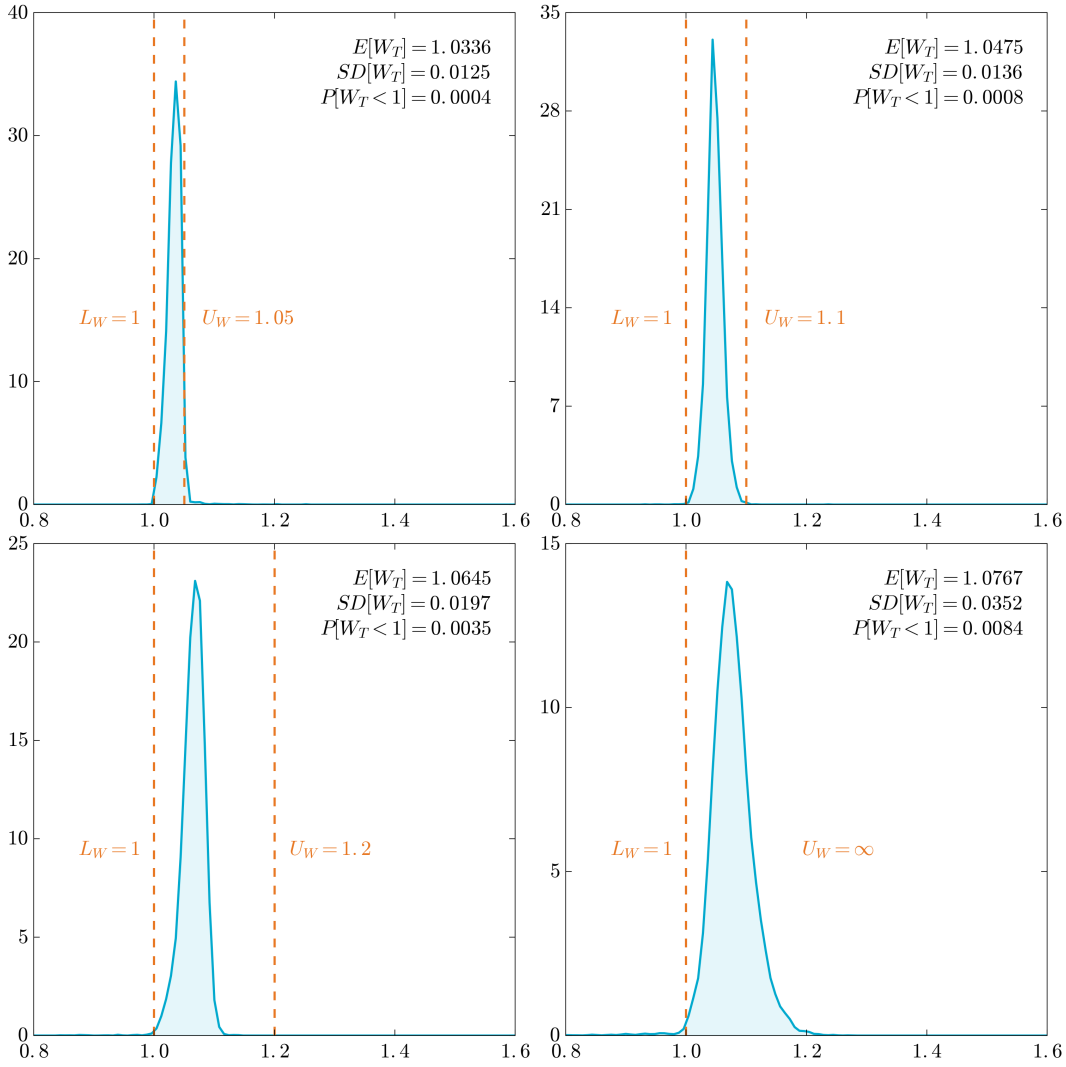


# 6 Extension 2: realized volatility bounds

This section describes an extension of the sharp target range strategy to incorporate a realized volatility target. The seminal paper of Markowitz (1952) on mean variance portfolio allocation ignited a vast literature on portfolio allocation techniques based on ex-ante estimated portfolio volatility. This abundant literature contrasts with the lack of techniques to account for the ex-post realized volatility in the investment decision process. In this section, we show how a target range on the terminal wealth level can be combined with a second target range on the realized volatility of the portfolio value. To avoid confusion, we make a clear distinction between the out-of-sample volatility of terminal portfolio value (the standard deviations reported on Figures 4.1, 4.2, 4.3, 4.4 and 5.2) and the realized volatility of the portfolio value over the trading period. We denote $\nu = (\nu_t)_{0 \leq t \leq T}$ as this realized volatility process,

defined by

$$\nu_t^2 := \frac{1}{t+1}\sum_{s=0}^{t}\left(W_s - \bar{W}_s\right)^2, \qquad \bar{W}_t := \frac{1}{t+1}\sum_{s=0}^{t} W_s\,. \tag{6.1}$$

The realized volatility process measures the variations of the investment portfolio value over time. While the sharp target range strategy (2.2) controls the volatility of terminal portfolio value, it does not explicitly control the ex-post realized volatility (6.1). Figure 6.1 shows an example of the distribution of terminal wealth (left picture) and the distribution of the corresponding realized volatility (right picture) for the target range $[L_W, U_W] = [1.0, 1.5]$. Even though the distribution of terminal wealth is well contained, the distribution of realized volatility is widely spread out.

Figure 6.1: Target return range strategy without realized volatility target range

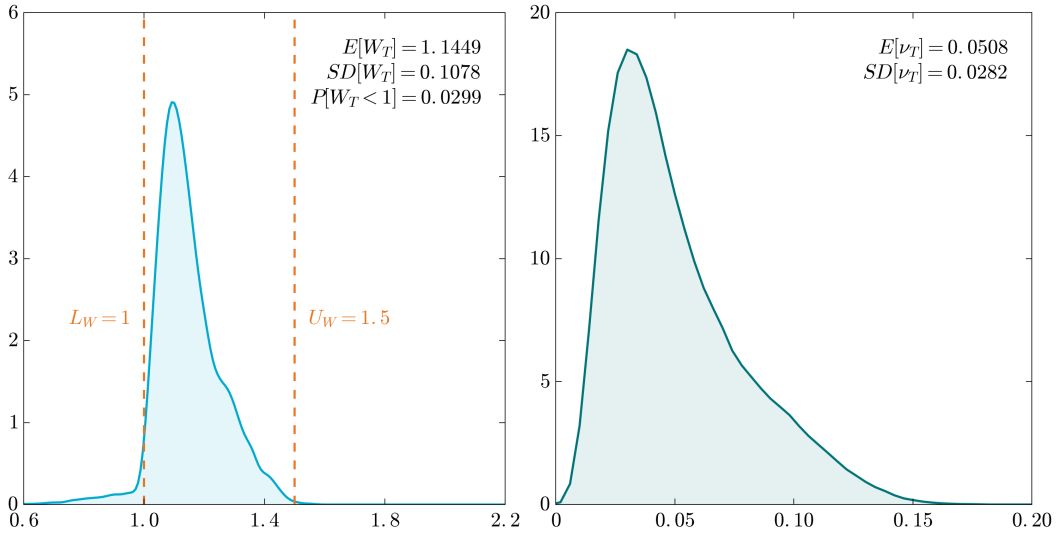


The purpose of this section is to introduce a new double target range strategy that combines the beneficial features of the target range (2.2) on returns with an additional target range on the realized volatility of the portfolio value. More specifically, we suggest a dynamic portfolio allocation aimed at maximizing the following objective function:

$$V_t = \max_{\{\alpha_s\}_{s=t}^{T-1}} \mathbb{E}[f_\nu(W_T, \nu_T)|\mathcal{F}_t] \tag{6.2}$$

where

$$f_\nu(w, v) := (w - L_W)\mathbb{1}\{L_W \le w \le U_W\,;\, L_\nu \le v \le U_\nu\} \tag{6.3}$$

where a new target range $[L_\nu, U_\nu]$ for realized volatility complements the target range $[L_W, U_W]$ for return. We effectively combine a sharp target range (2.2) for returns with a flat target range (5.1) for realized volatility. $L_\nu = 0$ is a natural lower target for realized volatility, while the realized volatility target $U_\nu$ can be set based on the investor's risk requirements.

We remark that the objective function (6.2) can be interpreted as an extension of the viability approach (De Lara and Doyen (2008), De Lara and Martinet (2009)), for which the joint probability of meeting a return target while staying below a risk target is maximized. In the context of portfolio allocation, the

viability approach would read

$$V_t = \max_{\{\alpha_s\}_{s=t}^{T-1}} \mathbb{P}[\{L_W \leq W_T\} \& \{\nu_T \leq U_\nu\} | \mathcal{F}_t] \,. \tag{6.4}$$

We finally remark that there is no restriction to extend (6.2)-(6.3) to include multiple target ranges for other risk measures, for example maximum drawdown, expected shortfall or beta risk.

## 6.1 Solution

In this subsection, we modify the algorithm introduced in Section (3) to solve the new objective function (6.2). The main difference with Section 3 is that there are now two endogenous state variables to account for: $W = (W_t)_{0 \leq t \leq T}$ and $\nu = (\nu_t)_{0 \leq t \leq T}$. We first describe the simulation procedure of the realized volatility. Given that the realized volatility (6.1) can be equivalently expressed as

$$\nu_t^2 = \bar{W}_t^{(2)} - \left(\bar{W}_t\right)^2 ,$$

where the average portfolio value $\bar{W}_t$ is given by equation (6.1), and the average squared portfolio value $\bar{W}_t^{(2)}$ is given by

$$\bar{W}_t^{(2)} \quad := \quad \frac{1}{t+1} \sum_{s=0}^{t} W_s^2,$$

the dynamics of $\nu$ can therefore be deduced from the dynamics of $\bar{W}_t$ and $\bar{W}_t^{(2)}$:

$$\begin{aligned} \bar{W}_t &= \frac{t}{t+1} \bar{W}_{t-1} + \frac{1}{t+1} W_t \\ \bar{W}_t^{(2)} &= \frac{t}{t+1} \bar{W}_{t-1}^{(2)} + \frac{1}{t+1} W_t^2. \end{aligned}$$

To compute the continuation values in the LSMC algorithm, we now perform a linear approximation of both final portfolio value $W_T$ and final realized volatility $\nu_T$

$$\begin{aligned} W_T &= \beta_W \cdot \psi(X_t^j) + \sigma_W \varepsilon_W, \ \varepsilon_W \sim \text{i.i.d. } \mathcal{N}(0,1), \\ \nu_T &= \beta_\nu \cdot \psi(X_t^j) + \sigma_\nu \varepsilon_\nu, \ \varepsilon_\nu \sim \text{i.i.d. } \mathcal{N}(0,1), \end{aligned}$$

at time $t$, as a function of all the state variables $X_t^j$ for each possible allocation $a^j \in \mathcal{C}_{\mathrm{d}}$, where the residuals $\varepsilon_W$ and $\varepsilon_\nu$ follow a joint standard normal distributions with correlation $\rho$. From there we can write

$$\begin{aligned} &\mathbb{E}[(W_T - L_W) \mathbb{1}\{L_W \leq W_T \leq U_W \,;\, L_\nu \leq \nu_T \leq U_\nu\} | X_t^j] \\ =&\mathbb{E}\left[\left(-\sigma_W L_W^j + \varepsilon_W\right) \mathbb{1}\{L_W^j \leq \varepsilon_W \leq U_W^j \,,\, L_\nu^j \leq \varepsilon_\nu \leq U_\nu^j\} | X_t^j\right] \end{aligned}$$

where

$$L_W^j := \frac{L_W - \beta_W \cdot \psi(X_t^j)}{\sigma_W}, \ U_W^j := \frac{U_W - \beta_W \cdot \psi(X_t^j)}{\sigma_W}, \ L_\nu^j := \frac{L_\nu - \beta_\nu \cdot \psi(X_t^j)}{\sigma_\nu}, \ U_\nu^j := \frac{U_\nu - \beta_\nu \cdot \psi(X_t^j)}{\sigma_\nu} \,. \tag{6.5}$$

Then,

$$\begin{aligned}&\mathbb{E}\left[\left(-\sigma_W L_W^j+\varepsilon_W\right)\mathbb{1}\{L_W^j\le\varepsilon_W\le U_W^j\,,\,L_\nu^j\le\varepsilon_\nu\le U_\nu^j\}|X_t^j\right]\\&=-\sigma_W L_W^j\mathrm{P}^j+\mathrm{E}^j\mathrm{P}^j\end{aligned}\tag{6.6}$$

where conditionally on $X_t^j$ (i.e. $L_W^j$, $U_W^j$, $L_\nu^j$ and $U_\nu^j$ being constant):

$$\begin{aligned}\mathrm{P}^j&:=\mathbb{P}\left[L_W^j\le\varepsilon_W\le U_W^j\,,\,L_\nu^j\le\varepsilon_\nu\le U_\nu^j\right]\\\mathrm{E}^j&:=\mathbb{E}\left[\varepsilon_W\mid L_W^j\le\varepsilon_W\le U_W^j\,,\,L_\nu^j\le\varepsilon_\nu\le U_\nu^j\right].\end{aligned}$$

We therefore need to compute $\mathrm{P}^j$ and $\mathrm{E}^j$. Firstly,

$$\mathrm{P}^j=\Phi(U_W,U_\nu;\rho)-\Phi(U_W,L_\nu;\rho)-\Phi(L_W,U_\nu;\rho)+\Phi(L_W,L_\nu;\rho)\tag{6.7}$$

where $\Phi(x,y;\rho)$ is the bivariate standard normal cumulative distribution function with correlation $\rho$.

Secondly, the distribution of $\varepsilon_W$ conditionally on $\{L_W^j\le\varepsilon_W\le U_W^j, L_\nu^j\le\varepsilon_\nu\le U_\nu^j\}$ is a bivariate truncated normal distribution, and its first moment $\mathrm{E}^j$ is derived in Equation (4) in Muthén (1990), with minor errors which are corrected in the Appendix of Kugiumtzis and Bora-Senta (2010). The corrected formula is

$$\begin{aligned}\mathrm{E}^j&=\mathbb{E}\left[\varepsilon_W\mid L_W^j\le\varepsilon_W\le U_W^j, L_\nu^j\le\varepsilon_\nu\le U_\nu^j\right]\\&=-\frac{\phi(U_W^j)}{P^j}\left(\Phi\left(\frac{U_\nu^j-\rho U_W^j}{\sqrt{1-\rho^2}}\right)-\Phi\left(\frac{L_\nu^j-\rho U_W^j}{\sqrt{1-\rho^2}}\right)\right)\\&\quad+\frac{\phi(L_W^j)}{P^j}\left(\Phi\left(\frac{U_\nu^j-\rho L_W^j}{\sqrt{1-\rho^2}}\right)-\Phi\left(\frac{L_\nu^j-\rho L_W^j}{\sqrt{1-\rho^2}}\right)\right)\\&\quad-\frac{\rho\phi(U_\nu^j)}{P^j}\left(\Phi\left(\frac{U_W^j-\rho U_\nu^j}{\sqrt{1-\rho^2}}\right)-\Phi\left(\frac{L_W^j-\rho U_\nu^j}{\sqrt{1-\rho^2}}\right)\right)\\&\quad+\frac{\rho\phi(L_\nu^j)}{P^j}\left(\Phi\left(\frac{U_W^j-\rho L_\nu^j}{\sqrt{1-\rho^2}}\right)-\Phi\left(\frac{L_W^j-\rho L_\nu^j}{\sqrt{1-\rho^2}}\right)\right).\end{aligned}\tag{6.8}$$

where as before $\phi(x)$ and $\Phi(x)$ represent the standard normal probability and cumulative distribution functions respectively.

To sum up, the continuation values for the double target range strategy are given by (6.6), which can be explicitly computed from equations (6.5), (6.7) and (6.8). The rest of the numerical procedure is then the same as in Section 3.

## 6.2 Out-of-sample distribution

Using the same numerical test case described in Section 4, we implement this new double target strategy and show out-of-sample distributions of both the final portfolio value and the realized volatility in Figure 6.2, where we fix the target return range to $[L_W,U_W]=[1.0,1.5]$, the lower realized volatility target $L_\nu=0$ and we test two upper realized volatility targets: $U_\nu=10\%$ (first row of Figure 6.2) and

$U_\nu = 5\%$ (second row of Figure 6.2).

Figure 6.2: Controlled realized volatility: out-of-sample distribution

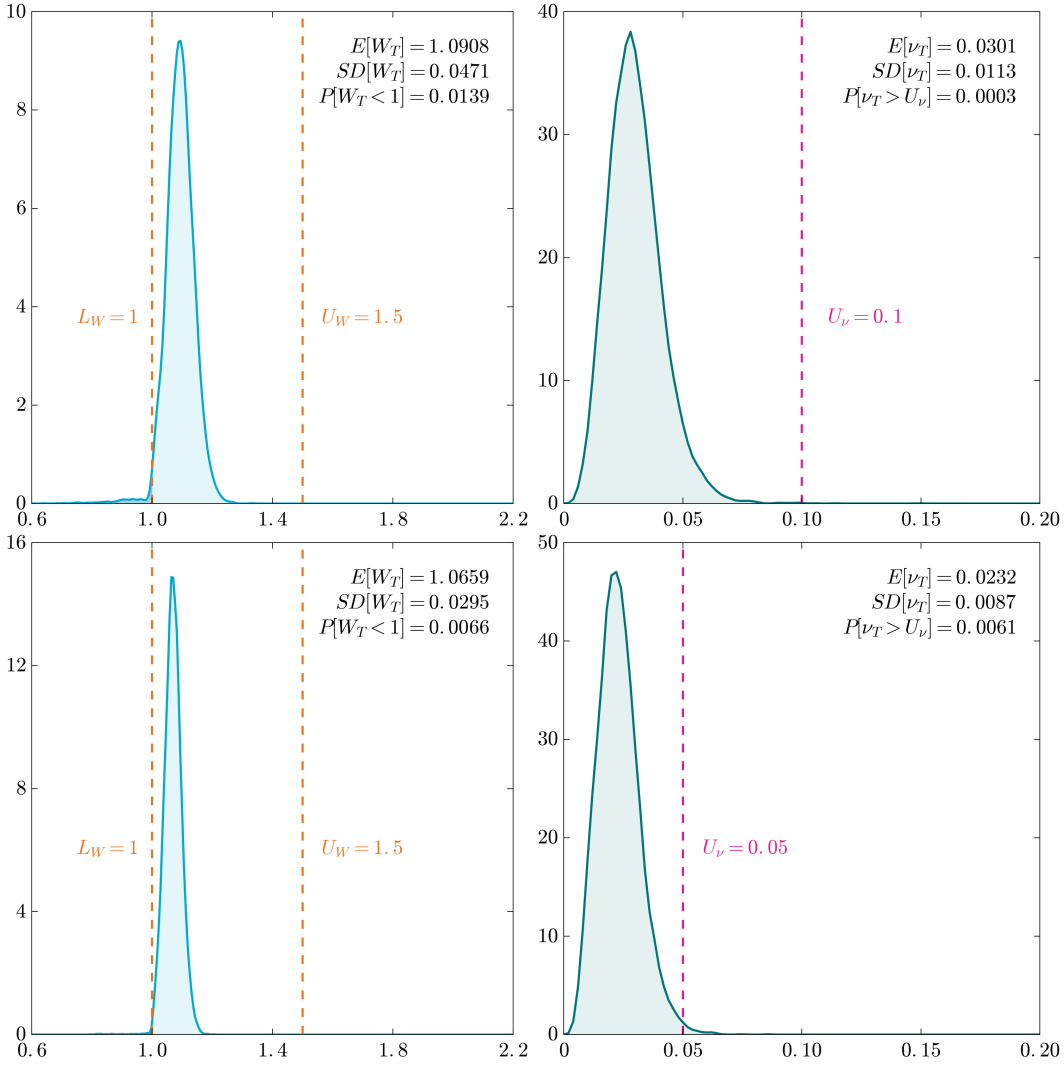


Comparing Figure 6.2 to Figure 6.1, the most striking observation is that the distribution of the realized volatility can be well targeted within the predetermined range such that less than 1% of the distribution lies outside the target. Moreover, standard deviation and downside risk of the final portfolio value are also drastically reduced. The price to pay for enforcing this realized volatility target is the decrease in expected return, as the range $[L_W, U_W] = [1.0, 1.5]$ generates 14.5% expected return with the sharp target range strategy ($U_\nu = +\infty$, Figure 6.1), down to only 9.1% and 6.7% for $U_\nu = 10\%$ and $U_\nu = 5\%$ respectively (Figure 6.2). This natural observation illustrates the unavoidable trade-off between risk and return. To analyze this phenomenon more precisely, we now perform a sensitivity analysis over the upper target of the realized volatility $U_\nu$.

Figure 6.3: Sensitivity analysis

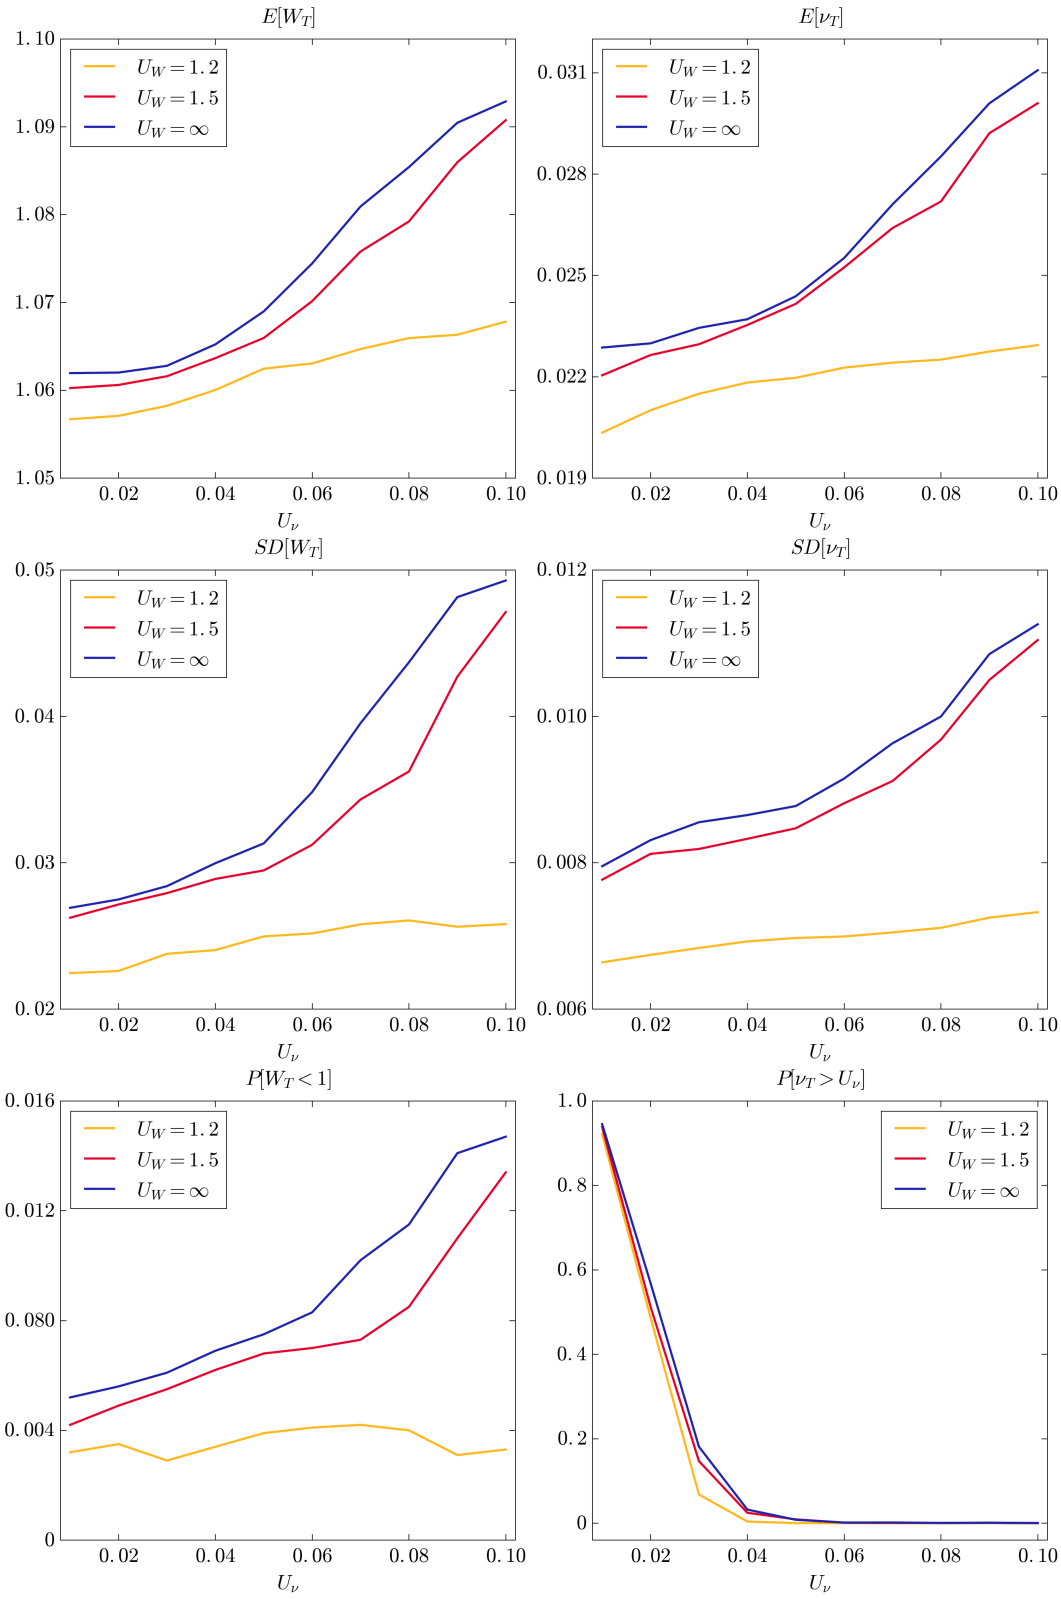

### 6.3 Sensitivity analysis

Figure 6.3 shows the sensitivity of the expectation, standard deviation and probability of achieving the return and realized volatility targets with respect to the upper realized volatility bound $U_\nu$. As expected, increasing $U_\nu$ increases the expectation and standard deviation of the final wealth and the realized volatility, increases the probability of not achieving the lower return target $\mathbb{P}(W_T < 1)$ and decreases the probability of not achieving the realized volatility target $\mathbb{P}(\nu_T > U_\nu)$. As shown, it is possible to reach relatively low levels of standard deviation and downside risk by decreasing $U_\nu$, which illustrates the flexibility provided by the new parameter $U_\nu$. However, these effects are modest when the upper return target is low, for example when $U_W = 1.2$, the distribution of the portfolio value is insensitive to the realized volatility's upper target. This is due to the fact that a low return target implicitly constrains the realized volatility target $\nu_T$, so that loosening the upper realized volatility bound $U_\nu$ does not bring enough additional flexibility to significantly change the portfolio performance.

Regarding the probability of not achieving the realized volatility target $\mathbb{P}(\nu_T > U_\nu)$, two main regimes stand out. This probability jumps high when $U_\nu < 4\%$, but is very small for $U_\nu > 4\%$ (negligible for $U_\nu > 5\%$), indicating that the minimum feasible realized volatility bound stands around 4%. This 4% minimum feasible volatility bound is the main finding of this subsection, and its knowledge can be beneficial in practice for conservative volatility-controlled funds.

## 7 Extension 3: stochastic benchmark target

So far, our target range strategies (2.2), (5.1) and (6.3) define the target returns as fixed, absolute returns. Yet, these strategies can also be defined relatively to a stochastic benchmark, be it stock index, inflation rate, exchange rate or interest rate. Most papers on classical target strategies aim at outperforming a stochastic benchmark (Franks (1992), Browne (1999a), Brogan and Stidham Jr. (2005), Gaivoronski et al. (2005) etc.), which is useful for tracking-based funds.

In this short section, we show how our target range strategies can easily be modified to accommodate relative target returns. Let $B = (B_t)_{0 \leq t \leq T}$ denote the stochastic benchmark of interest, and define the relative excess wealth as $W - B$. We can then modify the target range (2.2) as:

$$f_B(w, b) := (w - b - L_w)\mathbb{1}\{L_w \leq w - b \leq U_w\}\,. \tag{7.1}$$

Then, the new objective function is

$$V_t = \sup_{\{\alpha_s \in \mathcal{C}\}, s \in [t, T-1]} \mathbb{E}[f_B(W_T, B_T)|\mathcal{F}_t], t \in [0, T-1]\,. \tag{7.2}$$

By modeling the stochastic benchmark $B$ as an additional endogenous state variable, one can solve the dynamic portfolio allocation problem (7.2) using the same algorithm described in Section 3.

Using the same numerical test case described in Section 4 and using a passive, equal weight holding strategy as benchmark, we report the final distributions of excess return obtained with a sharp target range (Equation (7.1), Figure 7.1 top row), a flat target range (Figure 7.1 middle row), and a double return-and-realized-volatility range target (Figure 7.1 bottom row).

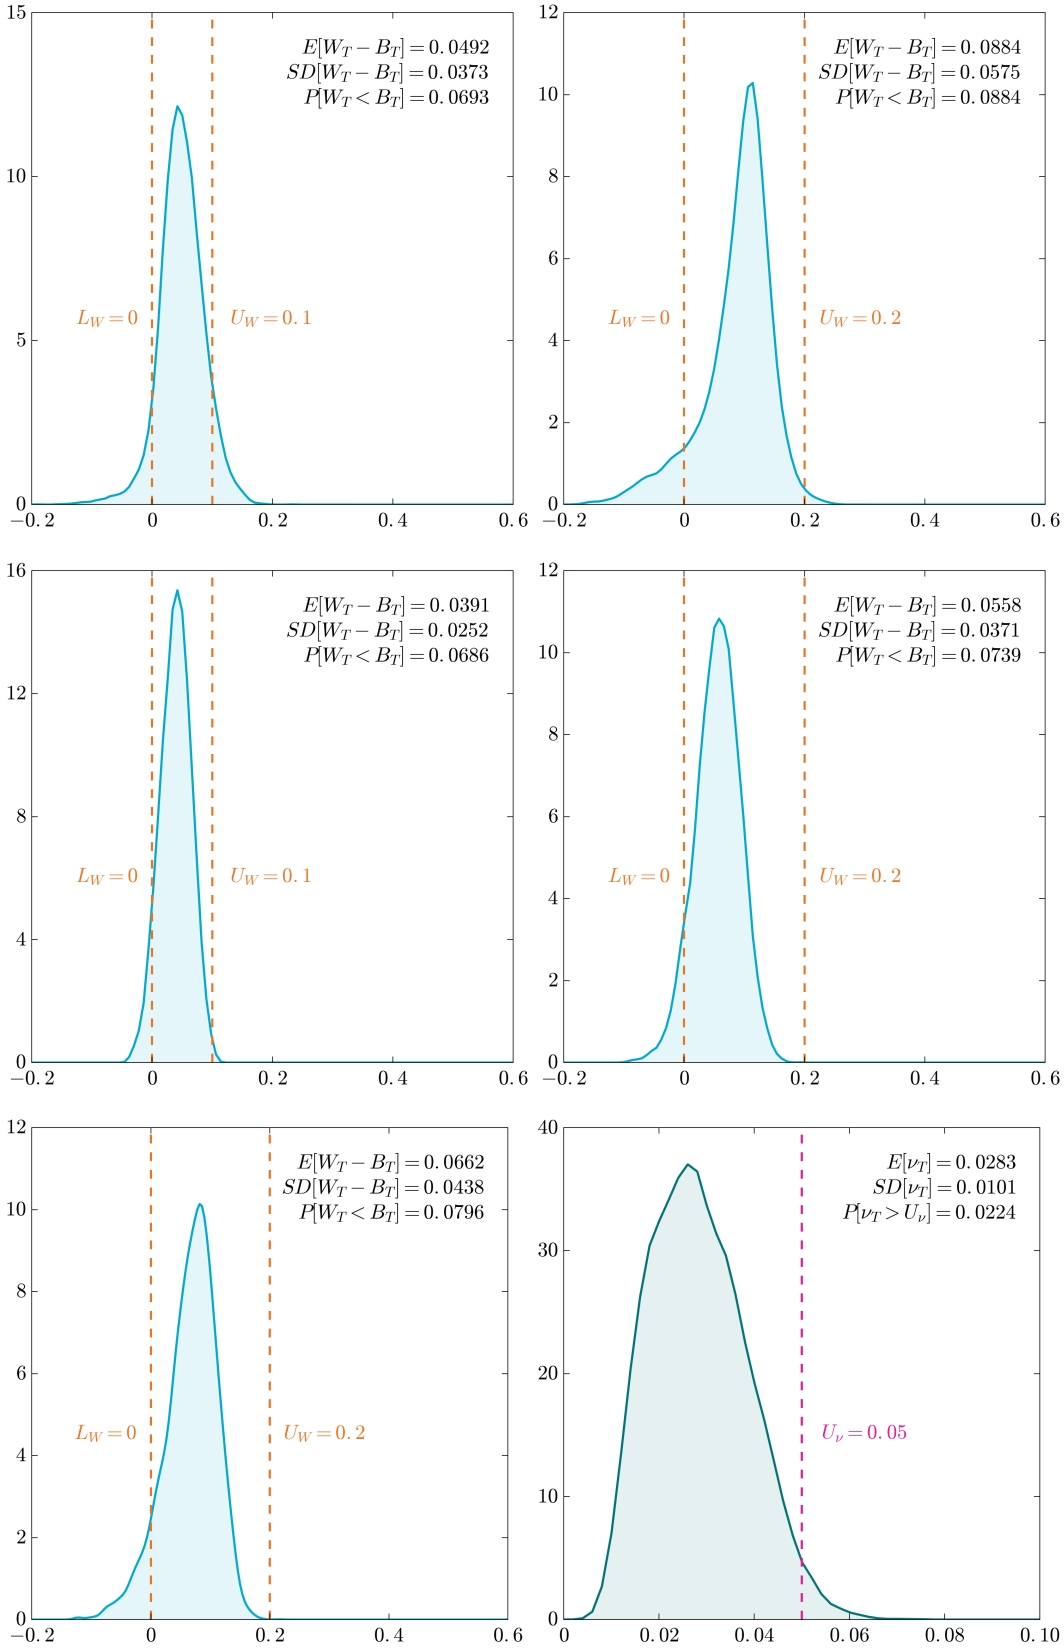


Figure 7.1: Relative return of target range (top row), probability maximizing (middle row), and double target range (bottom row) strategies

Figure 7.1 shows that our method can successfully deal with stochastic benchmarks as targets. The probability that the portfolio value ends outside the predefined ranges remains small (around $7\% - 8\%$ for the final excess return distributions, around 2% for the final realized excess volatility distribution), though higher than those provided by absolute range targets. This is because the passive equal weight benchmark already provides a significant positive expected return, therefore outperforming it requires to take more risk than that was needed previously for the absolute range target. Overall, relative target range is another useful application of our family of target range strategies.

# 8 Conclusion

This paper presents a new family of target strategies for portfolio selection problems. Our first strategy, dubbed sharp target range strategy, maximizes the expected portfolio value within a predefined target return range. This strategy manages the entire distribution of the terminal portfolio value: it simultaneously seeks high expected return, low standard deviation and limited downside risk. The simple parameterization accommodates various risk profiles and replicates classical hedge fund strategies (absolute return funds, directional funds, controlled volatility funds, etc.).

This sharp target range strategy can be seen as a better practical alternative to classical utility functions such as CRRA for several reasons. Firstly, unlike the abstract risk preference parameter of a standard utility function, the target range is expressed in the unit of investment return, making it simple, clear and intuitive. Secondly, for numerical approximations of multiperiod problems, the simplicity of our strategy avoids the instabilities resulting from the non-linearity of utility functions. Finally, our numerical tests show that the sharp target range strategy simply achieves a better efficient frontier and a better trade-off between return and downside risk than a classical constant relative risk aversion investment.

To demonstrate the effectiveness of these new strategies, we solve a dynamic portfolio selection problem in the presence of proportional transaction cost, using the LSMC algorithm. We show that the ex-post distributions are well targeted within the predefined ranges, replicating in a much simpler manner what could be expected from more involved constrained optimization methods. Three extensions are provided. Firstly, we consider a pure probability maximizing strategy, relevant for retirement funds and pensioners. Secondly, we incorporate an explicit target range on the realized volatility of the portfolio value. Our results show that a 3%-4% one-year monthly realized volatility can be achieved for a portfolio of cash, bond and stocks. Finally, we present a strategy that targets a range of excess returns over a stochastic benchmark, such as stock index or inflation rate.

Overall, this new family of investment strategies can be beneficial to the literature on portfolio allocation, whenever classical utility functions are used by default for lack of a better alternative, and also to portfolio management practitioners, who can use this new tool to achieve precise dynamic portfolio allocation targets.

# References


Bacon, C. (2013). *Practical risk-adjusted performance measurement.* Wiley.

Barberis, N. (2012). A model of casino gambling. *Management Science 58*(1), 35–51.

Brogan, A. J. and S. Stidham Jr. (2005). A note on separation in mean-lower-partial-moment portfolio optimization with fixed and moving targets. *IIE Transactions 37*(10), 901–906.

Browne, S. (1999a). Beating a moving target: Optimal portfolio strategies for outperforming a stochastic benchmark. *Finance and Stochastics 3*(3), 275–294.

Browne, S. (1999b). The risk and rewards of minimizing shortfall probability. *Journal of Portfolio Management 25*(4), 76–85.

Carriere, J. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics 19*(1), 19–30.

Davies, R., H. Kat, and S. Lu (2009). Fund of hedge funds portfolio selection: A multiple-objective approach. *Journal of Derivatives and Hedge Funds 15*(2), 91–115.

De Lara, M. and L. Doyen (2008). *Sustainable management of natural resources: mathematical models and methods*. Environmental science and engineering. Springer.

De Lara, M. and V. Martinet (2009). Multi-criteria dynamic decision under uncertainty: A stochastic viability analysis and an application to sustainable fishery management. *Mathematical Biosciences 217*(2), 118–124.

Franks, E. C. (1992). Targeting excess-of-benchmark returns. *Journal of Portfolio Management 18*(4), 6–12.

Gaivoronski, A. A., S. Krylov, and N. van der Wijst (2005). Optimal portfolio selection and dynamic benchmark tracking. *European Journal of Operational Research 163*(1), 115–131.

Garlappi, L. and G. Skoulakis (2009). Numerical solutions to dynamic portfolio problems: The case for value function iteration using Taylor approximation. *Computational Economics 33*, 193–207.

Hata, H., H. Nagai, and S.-J. Sheu (2010). Asymptotics of the probability minimizing a "down-side" risk. *Annals of Applied Probability 20*(1), 52–89.

Hodges, S. (1998). A generalization of the Sharpe ratio and its applications to valuation bounds and risk measures. Technical report, Financial Options Research Centre, University of Warwick.

Jorion, P. (2007). *Value at risk : the new benchmark for managing financial risk* (3rd ed.). New York: McGraw-Hill.

Kharroubi, I., N. Langrené, and H. Pham (2014). A numerical algorithm for fully nonlinear HJB equations: an approach by control randomization. *Monte Carlo Methods and Applications 20*(2), 145–165.

Klebaner, F., Z. Landsman, U. Makov, and J. Yao (2016). Optimal portfolios with downside risk. *Quantitative Finance*, 1–11.

Köksalan, M. and J. Wallenius (2012). Multiple criteria decision making: foundations and some approaches. *INFORMS Tutorial in Operations Research 8*, 171–183.

Konno, H., H. Shirakawa, and H. Yamazaki (1993). A mean-absolute deviation-skewness portfolio optimization model. *Annals of Operations Reseach 45*(1), 205–220.

Kugiumtzis, D. and E. Bora-Senta (2010). Normal correlation coefficient of non-normal variables using piece-wise linear approximation. *Computational Statistics 25*(4), 645–662.

Lai, T. (1991). Portfolio selection with skewness: A multiple-objective approach. *Review of Quantitative Finance and Accounting 1*(3), 293–305.

Longstaff, F. and E. Schwartz (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies 14*(1), 681–692.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance 7*(1), 77–91.

Milevsky, M. A., K. S. Moore, and V. R. Young (2006). Asset allocation and annuity-purchase strategies to minimize the probability of financial ruin. *Mathematical Finance 16*(4), 647–671.

Morton, D. P., E. Popova, and P. Ivilina (2006). Efficient fund of hedge funds construction under downside risk measures. *Journal of Banking and Finance 30*(2), 503–518.

Muthén, B. (1990). Moments of the censored and truncated bivariate normal distribution. *British Journal of Mathematical and Statistical Psychology 43*, 131–143.

Nagai, H. (2012). Downside risk minimization via a large deviation approach. *Annals of Applied Probability 22*(2), 608–669.

Pham, H. (2003). A large deviations approach to optimal long term investment. *Finance and Stochastics 7*(2), 169–195.

Rockafellar, R. (2007). Coherent approaches to risk in optimization under uncertainty. *INFORMS Tutorial in Operations Research 3*, 38–61.

Rockafellar, R. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Journal of Risk 2*(3), 21–42.

Stutzer, M. (2000). A portfolio performance index. *Financial Analysts Journal 56*(3), 52–61.

Teplá, L. (2001). Optimal investment with minimum performance constraints. *Journal of Economic Dynamics and Control 25*(10), 1629–1645.

Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica 26*(1), 24–36.

Tsitsiklis, J. and B. Van Roy (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks 12*(4), 694–703.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

Van Binsbergen, J. H. and M. Brandt (2007). Solving dynamic portfolio choice problems by recursing on optimized portfolio weights or on the value function? *Computational Economics 29*, 355–367.

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.

Williams, J. O. (1997). Maximizing the probability of achieving investment goals. *Journal of Portfolio Management 24*(1), 77–81.

Zakamouline, V. and S. Koekebakker (2009). Portfolio performance evaluation with generalized Sharpe ratios: beyond the mean and variance. *Journal of Banking and Finance 33*(7), 1242–1254.

Zhang, R., N. Langrené, Y. Tian, Z. Zhu, F. Klebaner, and K. Hamza (2016). Dynamic portfolio selection with permanent price impact: a flexible simulation-and-regression approach. Working paper, available at: https://arxiv.org/abs/1610.07694.